

Amendments to the Claims

1       Claim 1 (currently amended): A method of addressing a name space mismatch between content  
2       servers and content caching systems, comprising steps of:  
3                 caching one or more content elements in a content caching system, wherein the cached  
4       elements are identified by the content caching system using a first identification uniform resource  
5       locator format;  
6                 automatically learning mappings between a directory structure second identification  
7       format and the first identification uniform resource locator format, responsive to receiving the  
8       content elements for caching, wherein the received content elements are identifiable to the  
9       content caching system using the uniform resource locator format but also have associated  
10      therewith identifiers in the directory structure format, the mapping for each received content  
11      element storing its identifier in the uniform resource locator format and its associated identifier in  
12      the directory structure format; and  
13                 using the mappings to update when updating one or more selected cached content  
14       elements, responsive to receiving notifications of changed content elements that identify the  
15       changed content elements using the directory structure format, by requesting, by the content  
16       caching system, updated versions of the selected cached content elements from a content server,  
17       wherein each request identifies one of the selected cached content elements using the second  
18       identification identifier, in the uniform resource locator format, from the mapping that stores the  
19       identifier, in the directory structure format, from the received notification for that content  
20       element.

1       Claim 2 (currently amended): The method according to Claim 1, further comprising the step of  
2       using the mappings to invalidate one or more particular cached content elements by invoking an  
3       invalidation function for the content caching system, by the content caching system responsive to  
4       receiving notifications of invalid content elements that identify the particular content elements  
5       using the directory structure format, wherein each invocation identifies one of the particular  
6       cached content elements using the second naming identifier, in the uniform resource locator  
7       format, from the mapping that stores the identifier, in the directory structure format, from the  
8       received notification for that content element.

1       Claim 3 (currently amended): A method of addressing a name space mismatch between content  
2       servers and content caching systems, comprising steps of:

3              automatically learning mappings between a first identification directory structure format  
4       which identifies content elements accessible from one or more content servers and a second  
5       identification uniform resource locator format which identifies content elements cached by a  
6       caching system, when the caching system receives the content elements from the content servers  
7       for caching, wherein the received content elements are identifiable to the caching system using  
8       the uniform resource locator format and also have associated therewith corresponding identifiers  
9       in the directory structure format, the mapping for each received content element storing its  
10       identifier in the uniform resource locator format and its corresponding identifier in the directory  
11       structure format;

12              receiving, at the caching system, a notification pertaining to one of the cached content  
13       elements, wherein the notification identifies the content element using [[an]] its identifier in the

14 first identification directory structure format;  
15 consulting the learned mappings using the identifier in the first identification directory  
16 structure format from the notification, thereby determining [[a]] the corresponding identifier in  
17 identification of the cached content element, wherein the corresponding identification uses the  
18 second identification uniform resource locator format; and  
19 processing the received notification using the corresponding identification in the second  
20 identification uniform resource locator format.

1 Claim 4 (currently amended): The method according to Claim 3, wherein:  
2 the identifier in the first identification directory structure format identifies more than one  
3 of the cached content elements;  
4 the consulting step determines multiple corresponding identifications identifiers, in the  
5 uniform resource locator format, of the more than one identified cached content elements; and  
6 the processing step processes each of the identified cached content elements using the  
7 corresponding identifications identifiers in the second identification uniform resource locator  
8 format.

1 Claim 5 (original): The method according to Claim 3, wherein the notification is a content  
2 update notification.

1 Claim 6 (original): The method according to Claim 3, wherein the notification is a content  
2 invalidation notification.

1       Claim 7 (original): The method according to Claim 3, wherein the notification is received from a  
2       content management system.

1       Claim 8 (currently amended): The method according to Claim 3, wherein a plurality of  
2       notifications are received, the plurality comprising at least one content update notification and at  
3       least one content invalidation notification, and wherein the plurality of notifications use a  
4       consistent interface the directory structure format for conveying the notifications to the content  
5       caching system and also to the content servers.

1       Claim 9 (currently amended): The method according to Claim 3, wherein the mappings comprise  
2       one or more entries, and wherein the automatically learning step further comprises populating  
3       individual ones of the entries with the identification in the first identification format and the  
4       corresponding identification in the second identification format in response to receiving the  
5       content element to which the entry pertains operates when the content elements are received  
6       during cache miss processing.

1       Claim 10 (currently amended): The method according to Claim [9] 3, wherein the  
2       corresponding identifier populating step further comprises steps of:  
3       \_\_\_\_\_ extracting the identification in the first directory structure format is specified, for each of  
4       the received content elements, in [[from]] a response message with which carries the received  
5       content element is received, ; and

6 ————— using this extracted identification, along with the corresponding identification in the  
7 second identification format, as values of the populated entry.

1 Claim 11 (currently amended): A method of addressing a name space mismatch between content  
2 caching systems and content servers, comprising steps of:

3 sending a content request from a content caching system to a content server, wherein the  
4 content request identifies requested content using a first identification identifier in a first  
5 identification uniform resource locator format;

6 receiving a content response at the content caching system from the content server,  
7 wherein the content response provides the requested content, identifiable using the first identifier,  
8 and is augmented with a second identification identifier of the requested content [[using]] in a  
9 directory structure second identification format; and

10 creating an entry in a mapping, responsive to the receiving step, wherein the entry maps  
11 the second identification identifier in [[using]] the directory structure second identification  
12 format to the first identification identifier in [[using]] the first identification uniform resource  
13 locator format.

1 Claim 12 (original): The method according to Claim 11, wherein the content request signifies an  
2 ability to process the augmented content response.

1 Claim 13 (currently amended): The method according to Claim 12, wherein the ability is  
2 signified using a message header of the content request.

1       **Claim 14 (original): The method according to Claim 13, wherein the message header specifies**  
2       **that the caching system is content distribution aware.**

1       **Claim 15 (original): The method according to Claim 13, wherein the message header is a**  
2       **Hypertext Transfer Protocol ("HTTP") header.**

1       **Claim 16 (original): The method according to Claim 12, wherein the ability is signified using**  
2       **syntax of a markup language.**

1       **Claim 17 (original): The method according to Claim 16, wherein the markup language is**  
2       **Hypertext Markup Language ("HTML").**

1       **Claim 18 (original): The method according to Claim 17, wherein the syntax comprises a**  
2       **"META" tag using an "HTTP-EQUIV" attribute syntax.**

1       **Claim 19 (original): The method according to Claim 17, wherein the syntax comprises a**  
2       **"META" tag using a "NAME" attribute syntax.**

1       **Claim 20 (original): The method according to Claim 17, wherein the syntax comprises a**  
2       **specially-denoted comment.**

1       Claim 21 (original): The method according to Claim 16, wherein the markup language is XML  
2       ("Extensible Markup Language").

1       Claim 22 (original): The method according to Claim 12, wherein the ability is signified using a  
2       cookie on the content request.

1       Claim 23 (currently amended): The method according to Claim 11, wherein the augmented  
2       content response provides the second identification identifier using a header.

1       Claim 24 (original): The method according to Claim 23, wherein the header comprises an  
2       extended cache control header.

1       Claim 25 (original): The method according to Claim 23, wherein the header comprises a file  
2       identification header.

1       Claim 26 (original): The method according to Claim 25, wherein the file identification header is  
2       a Hypertext Transfer Protocol ("HTTP") header.

1       Claim 27 (currently amended): The method according to Claim 11, wherein the augmented  
2       content response provides the second identification identifier using syntax of a markup language.

1       Claim 28 (original): The method according to Claim 27, wherein the markup language is

2 Hypertext Markup Language ("HTML").

1 Claim 29 (original): The method according to Claim 28, wherein the syntax comprises a  
2 "META" tag using an "HTTP-EQUIV" attribute syntax.

1 Claim 30 (original): The method according to Claim 28, wherein the syntax comprises a  
2 "META" tag using a "NAME" attribute syntax.

1 Claim 31 (original): The method according to Claim 28, wherein the syntax comprises a  
2 specially-denoted comment.

1 Claim 32 (original): The method according to Claim 27, wherein the markup language is XML  
2 ("Extensible Markup Language").

1 Claim 33 (currently amended): The method according to Claim 11, wherein the augmented  
2 content response provides the second-identification identifier using a cookie on the content  
3 request.

1 Claim 34 (original): The method according to Claim 11, further comprising the step of caching  
2 the received content at the content caching system, responsive to the receiving step.

1 Claim 35 (currently amended): The method according to Claim 34, further comprising the steps

2 of:

3 receiving an update notification for a selected previously-cached content at the content  
4 caching system, wherein the update notification identifies the selected previously-cached content  
5 using a selected identification identifier thereof in the second identification directory structure  
6 format;

7 consulting the mapping, using the selected identification identifier as the second  
8 identification in the second identification format, to retrieve the first identification identifier in  
9 the first identification uniform resource locator format; and

10 using the retrieved first identification identifier to request an update of the selected  
11 previously-cached content.

1 Claim 36 (currently amended): The method according to Claim 35, wherein the using step  
2 further comprises steps of:

3 sending a new content request to the content server, wherein the new content request  
4 identifies the requested content using the retrieved first identification identifier;

5 receiving a new content response from the content server, wherein the new content  
6 response provides the requested updated content; and

7 caching the received updated content.

1 Claim 37 (currently amended): The method according to Claim 34, further comprising the steps  
2 of:

3 receiving an invalidation notification for a selected previously-cached content at the

4       content caching system, wherein the invalidation notification identifies the selected previously-  
5       cached content using a selected identification identifier in the second identification directory  
6       structure format;

7               consulting the mapping, using the selected identification identifier as the second  
8       identification in the second identification format, to retrieve the first identification identifier in  
9       the first identification uniform resource locator format; and

10              using the retrieved first identification identifier to request an invalidation of the selected  
11       previously-cached content.

Claims 38 - 40 (canceled)

1       Claim 41 (currently amended): The method according to Claim 11, wherein the second  
2       identification identifier is a file path and a file name.

1       Claim 42 (currently amended): The method according to Claim 11, wherein the second  
2       identification identifier is a substituted file path and file name.

1       Claim 43 (original): The method according to Claim 42, wherein the substituted file path and  
2       file name is an encrypted file path and file name.

1       Claim 44 (currently amended): The method according to Claim 11, further comprising steps of:  
2              receiving, at the content server, the sent content request;

3 obtaining, by the content server, the requested content;  
4 preparing a response message, by the content server, to transmit the obtained content,  
5 wherein the response message is augmented with the second identification identifier in the  
6 directory structure second identification format; and  
7 sending the prepared response message, including the obtained content, from the content  
8 server to the content caching system.

1 Claim 45 (currently amended): A system for addressing a name space mismatch between content  
2 servers and content caching systems, comprising:

3 means for automatically learning mappings between a first identification directory  
4 structure format which identifies content elements accessible from one or more content servers  
5 and a second identification uniform resource locator format which identifies content elements  
6 cached by a caching system, when the caching system receives the content elements from the  
7 content servers for caching, wherein the received content elements are identifiable to the caching  
8 system using the uniform resource locator format and also have associated therewith  
9 corresponding identifiers in the directory structure format, the mapping for each received content  
10 element storing its identifier in the uniform resource locator format and its corresponding  
11 identifier in the directory structure format;

12 means for receiving at the caching system, a notification pertaining to one of the cached  
13 content elements, wherein the notification identifies the content element using [[an]] its identifier  
14 in the first identification directory structure format;

15 means for consulting the learned mappings using the identifier in the first identification

16       directory structure format from the notification, thereby determining [[a]] the corresponding  
17       identifier in identification of the cached content element, wherein the corresponding  
18       identification uses the second identification uniform resource locator format; and  
19               means for processing the received notification using the corresponding identification in  
20       the second identification uniform resource locator format.

1       Claim 46 (currently amended): A system for addressing a name space mismatch between content  
2       caching systems and content servers, comprising:

3               means for sending a content request from a content caching system to a content server,  
4       wherein the content request identifies requested content using a first identification identifier in a  
5       first identification uniform resource locator format;

6               means for receiving a content response at the content caching system from the content  
7       server, wherein the content response provides the requested content, identifiable using the first  
8       identifier, and is augmented with a second identification identifier of the requested content  
9       [[using]] in a directory structure second identification format; and

10               means for creating an entry in a mapping, responsive to the means for receiving, wherein  
11       the entry maps the second identification identifier in [[using]] the directory structure second  
12       identification format to the first identification identifier in [[using]] the first identification  
13       uniform resource locator format.

1       Claim 47 (currently amended): A computer program product for addressing a name space  
2       mismatch between content servers and content caching systems, the computer program product

3 embodied on one or more computer-readable media and comprising:

4 computer readable program code means for automatically learning mappings between a  
5 first identification directory structure format which identifies content elements accessible from  
6 one or more content servers and a second identification uniform resource locator format which  
7 identifies content elements cached by a caching system, when the caching system receives the  
8 content elements from the content servers for caching, wherein the received content elements are  
9 identifiable to the caching system using the uniform resource locator format and also have  
10 associated therewith corresponding identifiers in the directory structure format, the mapping for  
11 each received content element storing its identifier in the uniform resource locator format and its  
12 corresponding identifier in the directory structure format;

13 computer readable program code means for receiving, at the caching system, a  
14 notification pertaining to one of the cached content elements, wherein the notification identifies  
15 the content element using [[an]] its identifier in the first identification directory structure format;

16 computer readable program code means for consulting the learned mappings using the  
17 identifier in the first identification directory structure format from the notification, thereby  
18 determining [[a]] the corresponding identifier in identification of the cached content element,  
19 wherein the corresponding identification uses the second identification uniform resource locator  
20 format; and

21 computer readable program code means for processing the received notification using the  
22 corresponding identification in the second identification uniform resource locator format.

1 Claim 48 (currently amended): A computer program product for addressing a name space

2 mismatch between content servers and content caching systems, the computer program product  
3 embodied on one or more computer-readable media and comprising, comprising:

4 computer readable program code means for sending a content request from a content  
5 caching system to a content server, wherein the content request identifies requested content using  
6 a first identification identifier in a first identification uniform resource locator format;

7 computer readable program code means for receiving a content response at the content  
8 caching system from the content server, wherein the content response provides the requested  
9 content, identifiable using the first identifier, and is augmented with a second identification  
10 identifier of the requested content [[using]] in a directory structure second identification format;

11 and

12 computer readable program code means for creating an entry in a mapping, responsive to  
13 the computer readable program code means for receiving, wherein the entry maps the second  
14 identification identifier in [[using]] the directory structure second identification format to the first  
15 identification identifier in [[using]] the first identification uniform resource locator format.